

Requested Patent: WO9850838A2

Title:

A METHOD AND SYSTEM FOR PROVIDING ON-LINE INTERACTIVITY OVER A
SERVER-CLIENT NETWORK ;

Abstracted Patent: WO9850838 ;

Publication Date: 1998-11-12 ;

Inventor(s):

GUELDALIA ISAAC DAVID (IL); HASHKES JONATHAN (IL); GUEDALIA JACOB
LEON (IL) ;

Applicant(s):

LIVE PICTURE INC (US); GUELDALIA ISAAC DAVID (IL); HASHKES JONATHAN
(IL); GUEDALIA JACOB LEON (IL) ;

Application Number: WO1998IL00174 19980409 ;

Priority Number(s): IL19980123479 19980226; US19970850690 19970502 ;

IPC Classification: G06F ;

Equivalents:

AU6851198, CA2287904, EP0992003, A2, A3, JP2000511321T, JP2003099358,
JP3457334B2, US6535878 ;

ABSTRACT:

Apparatus and a method for communicating media over a network including encoding the media into a server database at a server, downloading from the server database to a client database generally only those portions of the media which are necessary to satisfy user requests and in response to a user request for a given item of media, determining whether the media is present in the client database, and if not, automatically downloading those portions of the media which are necessary to supply the user with the given item of media from the server database. An image server is provided including at least one processor operative to supply portions of image data to clients in response to multiple requests therefrom, and thread management software operating the at least one processor by causing it to process the requests using at least one of a plurality of threads, the thread management software being characterized in that it initiates a new thread when an existing thread has exceeded a predetermined metric of busyness.

BEST AVAILABLE COPY

PCTWORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F	A2	(11) International Publication Number: WO 98/50838 (43) International Publication Date: 12 November 1998 (12.11.98)
(21) International Application Number: PCT/IL98/00174 (22) International Filing Date: 9 April 1998 (09.04.98) (30) Priority Data: 08/850,690 2 May 1997 (02.05.97) US 123479 26 February 1998 (26.02.98) IL (71) Applicant (for all designated States except US): LIVE PICTURE, INC. [US/US]; 910 East Hamilton Avenue, Campbell, CA 95008 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): GUELDALIA, Isaac, David [IL/IL]; Shimon Street 80, 99000 Bet Shemesh (IL). HASHKES, Jonathan [IL/IL]; Shmaryahu Levin 20, 96664 Jerusalem (IL). GUEDALIA, Jacob, Leon [IL/IL]; Haporzim Street 24, 92541 Jerusalem (IL). (74) Agents: SANFORD, T., Colb et al.; Sanford T. Colb & Co., P.O. Box 2273, 76122 Rehovot (IL).		(81) Designated States: AL, AM, AT, AT (Utility model), AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, CZ (Utility model), DE, DE (Utility model), DK, DK (Utility model), EE, EE (Utility model), ES, FI, FI (Utility model), GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SK (Utility model), SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published <i>Without international search report and to be republished upon receipt of that report.</i>
(54) Title: A METHOD AND SYSTEM FOR PROVIDING ON-LINE INTERACTIVITY OVER A SERVER-CLIENT NETWORK		
(57) Abstract <p>Apparatus and a method for communicating media over a network including encoding the media into a server database at a server, downloading from the server database to a client database generally only those portions of the media which are necessary to satisfy user requests and in response to a user request for a given item of media, determining whether the media is present in the client database, and if not, automatically downloading those portions of the media which are necessary to supply the user with the given item of media from the server database. An image server is provided including at least one processor operative to supply portions of image data to clients in response to multiple requests therefrom, and thread management software operating the at least one processor by causing it to process the requests using at least one of a plurality of threads, the thread management software being characterized in that it initiates a new thread when an existing thread has exceeded a predetermined metric of busyness.</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

A Method and System for Providing On-Line Interactivity over a Server-Client Network

FIELD OF THE INVENTION

The present invention relates to server-client networks generally, and more specifically to server thread management which enables on-line access of interactive media over the Internet.

BACKGROUND OF THE INVENTION

Sharing of information has long been a goal of modern civilization. Information must be represented in some form in order that it can be transferred between people. Humans are gifted with five sensory modes of perception. Visual perception has been given the lion's share of the brain. It is noted by Dr. Judith Guedalia, of the Neuro-Psychological Department of Shaare Zedek Medical Center in Jerusalem, Israel, that although the famous Helen Keller, who was both deaf and blind, was quoted as saying that she would prefer hearing over sight if she had to choose, this may have been based on psychological reasons, since group activity is better enabled with hearing.

Visual stimulation is perhaps the most efficient means for people to acquire information. The cliché "a picture is worth a thousand words" is well-founded. The ratio of brains cells devoted to sight versus those devoted to sound is roughly 100:1. From the days of cave carving to the printing press, visual information has been the primary medium for communication.

Recently information has taken on a digital form, which has enabled electronic transmission of information. The most notable example is the dissemination of information through the World Wide Web (referred to simply as the "Web"), which is a collection of millions of computers interconnected by electronic communication devices.

Various techniques are known for providing on-line access of interactive media over the Web.

The following U.S. Patents have been found in a U.S. Patent Search and are believed to be generally relevant to the field of the invention:

4,897,867	1/90 Foster et al.
5,119,188	6/92 McCalley et al.
5,122,873	6/92 Golin
5,195,092	3/93 Wilson et al.
5,220,420	6/93 Hoarty et al.
5,236,199	8/93 Thompson, Jr.
5,251,209	10/93 Jurkevich et al.
5,265,248	11/93 Moullos et al.
5,283,819	1/94 Glick et al.
5,325,423	6/94 Lewis
5,351,276	9/94 Doll, Jr. et al.
5,363,482	11/94 Victor et al.
5,420,572	5/95 Dolin, Jr. et al.
5,420,801	5/95 Dockter et al.
5,438,658	8/95 Fitzpatrick et al.
5,487,167	1/96 Dinallo et al.
5,495,576	2/96 Ritchey
5,508,940	4/96 Rossmere et al.
5,519,435	5/96 Anderson
5,553,221	9/96 Reimer et al.
5,553,222	9/96 Milne et al.
5,557,538	9/96 Retter et al.
5,561,791	10/96 Mendelson et al.
5,564,001	10/96 Lewis
5,577,180	11/96 Reed
5,577,258	11/96 Cruz et al.
5,581,783	12/96 Ohashi

Transportation of digital signals between computers is plagued by a bandwidth problem, where a limited bandwidth creates bottlenecks in the transmission of electronic signals. Fortunately, textual information can be represented efficiently electronically, and this has enabled the Web publishing industry to flourish. Unfortunately, however, image information is difficult to represent compactly in electronic form.

A server or host computer (hereinafter referred to as a "server") is used for placing information where it is available for access by multiple users. The Web has enabled convenient publishing of information from a server to client computers (hereinafter referred to as "clients") that

are controlled by users requesting information.

When using various media such as video, audio, text and images, a user generally retrieves the media from a server connected via a network to many clients. The server downloads the media to the network and transmits it to the client computer at the user's request.

There are two basic limitations involved in such data retrieval: delay between the time that a user requests the data and the time when the server downloads it to the network, and bandwidth limitations on data throughput and rate of data transmission.

One example of such a server-client network is a network connecting Web servers and users' personal computers. Such networks are installed in order to facilitate convenient data transmission between users, and data distribution from the server to the users' computers.

Known network applications involve streaming data from a server to a client computer. "Streaming" refers to serial or parallel transmission of digital data between two computers, by transmitting sequences of bit packets. For example, installation executables on a network server stream files to a client computer performing the installation. Servers with large amounts of memory are used to archive digital movies, which are streamed to a client computer for viewing upon demand. Digital video is broadcast from cable stations to subscribers using streaming. Web browsers, such as Netscape and Microsoft Explorer, are used to stream data from a server on the Web to a client.

Web sites can contain enormous databases, such as phone directories for all of the cities in the U.S., photographs from art galleries and museums around the world, voluminous encyclopedias, and even copies of all patents ever issued by the U.S. Patent & Trademark Office. Users of the Web can search these databases and then request the server to download specific information. This request initiates a streaming event.

Accessing information over the Web is typically done using Web browsers. A Web browser is software that resides on a client computer and communicates with servers via established protocols. Information transmitted by servers is converted to visual displays on the client computer by means of a Web browser.

Internet protocols enable client-server communication over

the Web. These protocols include low level protocols, such as Transaction Control Protocol/Internet Protocol (TCP/IP), and higher level protocols such as Hypertext Transfer Protocol (HTTP). A general reference on Internet protocols may be accessed on:

<http://www.w3.org/Protocols>,

or

<http://www.cis.ohio-state.edu/tbin/rfc/arpa-internet-protocols.html>.

Another useful reference is: Hethmon, Paul S., *Illustrated Guide to HTTP*, Manning Publications Co., Greenwich, CT, 1997.

HTTP servers provide a way for remote clients to access data on servers. HTTP provides a method for making simple requests from a client to a server. Client requests can take the form of GET requests or POST requests in HTTP 1.0. Typically, in a GET or POST request the client specifies a file to be delivered, and through HTTP headers the server can specify what is being delivered. The most pervasive file format used on the Web is HyperText Markup Language (HTML). A reference on HTML may be accessed at:

<http://204.57.196.12/reference/htmlspec2.0>.

HTML files are typically relatively small, i.e. less than 100 Kbytes.

HTTP/1.0 specifies that a communication between a server and client proceeds as follows: A client's request is initiated with a header which is terminated by a double carriage return linefeed. This is followed by the body of the request which is similarly terminated by a double carriage return linefeed. The server responds with an HTTP header terminated with a carriage return linefeed and then sends the body of the response. The response is terminated when a socket is closed. This normally occurs when the server has finished returning data to the client, and the server closes the socket.

Server performance is generally inversely proportional to the quantity of data being served per unit time. The task of delivering a file from a server to a client is typically not computationally expensive. This task includes reading the file from the server's peripherals, e.g. a disk drive, and transmitting the data from the file in the specified protocol, e.g. TCP/IP. TCP/IP transmits data in units of "packets." The time it takes for a client to retrieve a file depends upon the number of packets transmitted by the server.

One problem shared by virtually all server applications is the need to process many client requests simultaneously. Efficient methods of processing multiple client requests utilize the server's resources fully. Many computers, even though they are limited by having only a single Central Processing Unit (CPU), can perform specific multiple independent tasks simultaneously, such as reading data from the disk and transmitting data over a network connection. This is enabled by buffers designed to speed up the dialogue with the peripherals, such as network cards, disks and video display cards.

Modern operating systems support multi-tasking, which is the ability to run many separate applications at the same time. A single software program can take advantage of multi-tasking by creating multiple concurrent "threads." Each thread simulates a separate application. Thus an HTTP server, for example, can use multiple threads to optimize its performance in responding to concurrent requests. Each request can be processed in a separate thread, and while one request is being processed in the CPU, a second request can be transmitted through the network hardware. If only one thread is used, then although the network hardware is buffered, the processing of the second request can become blocked – because the single thread waits for the network to finish sending.

Another advantage of using multiple threads to handle multiple requests is fairness. For example, suppose two clients simultaneously request information, one requests a large quantity of data while another requests a small amount. If their requests were to be handled within a single thread, the second client would have to wait until the completion of the first request. Multiple threads enable both clients to receive data concurrently.

The use of multiple threads in server applications works best when each request is independent of the other, in the sense that, for the most part, they do not share the same host computer resources. However, some sharing is inevitable. When threads share resources, this may produce bottlenecks in the transmission of data, CPU usage and access to peripherals. For example, using multiple threads creates problems such as page faults. Page faults occur when a thread requests a page of memory which is not readily available to the operating system. Thus server

applications try to limit the number of concurrent threads, while optimizing performance and fairness.

One popular method for limiting the number of concurrent threads, implemented in the servers of two leading commercial server vendors -- Microsoft Internet Information Services (IIS) and Netscape Enterprise, is as follows: A maximum number of threads is pre-set. Multiple threads are created as requests are received. In these servers, thread initiation is determined by client requests. These threads form a thread pool. The threads can be in either of two states -- a wait state, in which they are idle and wait for a request to process, or a busy state, in which they are processing a request. At a time that the maximum limit of allowable concurrent requests/threads is reached and they are all busy, subsequent requests are queued pending availability of a free thread.

Added functionality for server-side applications may be provided via a secondary protocol, such as Common Gateway Interface (CGI). A reference on CGI may be accessed at:

<http://www.pricecostco.com/exchange/irf/cgi-spec.html>.

Reference is also made to Dwight, J., Erwin, Michael and Niles, Robert, *Using CGI, Second Edition*, Que Corporation, Indianapolis, IN, 1997.

Unlike file retrieval governed by HTML, added functionality may involve intensive processing. The introduction of many server-side applications introduces severe performance problems. As mentioned above, multiple threads work best when each request is independent of each other. Otherwise, if two requests use a common resource, e.g., the same database, a situation may occur in which one request needs to wait until processing of the other request is complete and the resource is freed up. Currently webmasters, who are administrators of servers on the Web, manually tune a number of parameters to increase performance of their Web servers. A description of fine-tuning thread administration may be accessed at:

<http://www.asia.microsoft.com/isn/techcenter/tuningIIS.htm>.

In the above-referenced description there appears a discussion of how to tune two parameters, `ProcessorThreadMax` and `RequestQueueMax`. `ProcessorThreadMax` sets a maximum limit on the number of concurrent threads per processor, and `RequestQueueMax` sets a

maximum limit on the size of the request queue. Reference is also made to *Microsoft Windows NT Server Resource Kit Version 4.0, Supplement One*, Microsoft Press, Redmond, WA, 1997, Chap. 6 entitled "Preventing Processor Bottlenecks." This reference also discusses limiting the maximum number of threads and limiting the maximum number of connections. Quoting from page 136, "*If you cannot upgrade or add processors, consider reducing the maximum number of connections that each IIS service accepts. Limiting connections might result in connections being blocked or rejected, but it helps ensure that accepted connections are processed promptly. ... In extreme cases of very active or underused processors, you might want to adjust the maximum number of threads in the Inetinfo process.*"

It may thus be appreciated that while a simple thread model performs adequately for an HTTP server that serves HTML documents, this same model may fail for CPU intensive server-side scripts.

Users desiring to access image information over the Web run up against the bandwidth problem. Whereas even complex text information, such as multiple fonts and scales, can be represented with a few dozen bytes of HTML, images are usually orders of magnitude larger. Images are typically represented in one of two standard formats, JPEG or GIF. Currently, to transmit a small photograph over the Web, such as a 2" x 3" photograph scanned in at 300 dots per inch (dpi) and compressed using JPEG, takes more than two minutes over a typical 1 KByte per second connection. This makes the viewing of quality images, such as a small museum painting 10" x 10" scanned in at 300 dpi and JPEG compressed, very frustrating.

A recently developed protocol, the Internet Imaging Protocol (IIP), was designed specifically for mitigating the bandwidth problem. It exploits the fact that the client cannot view more than a computer screen of image data at any one time. Even if the full image is enormous, such as 15,000 x 15,000 pixels, the client never views more than the screen resolution, usually less than 1,024 x 1,024 pixels, at any given time. Thus it is unnecessary to transmit more than a screen-ful of data, for any specific view of the image. IIP specifies a method for a client to request portions of an image at a specific resolution. A reference for IIP is the

document "Internet Imaging Protocol," ©1997 Hewlett Packard Company, Live Picture, Inc., and Eastman Kodak Company, the contents of which are hereby incorporated by reference.

A server with server-side software that supports IIP is referred to as an "image server." Currently there are two popularly accepted ways to request image data from an image server using IIP; namely, using server-side processing of the request primarily, or using client-side processing primarily. Client-side processing is not relevant to the present invention, and will not be described herein.

To illustrate server-side processing, suppose a client with a viewing window of 640 x 480 pixels desires to view an image whose full resolution is 15,000 x 15,000 pixels. The client is unable to view the full image at its original resolution, and can either view the complete image at a low resolution, or view only a portion of the image at a high resolution. Usually the user prefers to begin with an initial view showing the full image at a low resolution, and then to interactively navigate by zooming, i.e., increasing the resolution while decreasing the "field of view," or the portion of the image being viewed, and panning, i.e. translating the current view.

Under IIP, the full image at a 640 x 480 pixel resolution for an initial view can be requested using the following set of IIP commands:

```
fif=<image name>&wid=640&hei=480&cvt=jpeg
```

This request specifies the desired image by means of the *fif* command, and specifies the width and height of the client viewing window by means of the *wid* and *hei* commands, respectively. The last command, *cvt*, specifies the format of the image to be returned. As mentioned hereinabove, JPEG is supported by most browsers.

For the image server to process the above IIP request, the server must analyze the original image and generate an image matching the requested specifications, specifically the desired portion and dimensions. The analysis and generation are usually computationally expensive. In the example under discussion, a 15,000 x 15,000 pixel image would have to be re-sized, requiring approximately 675 MBytes to process.

To facilitate this processing, images can be stored in a pre-computed multi-resolution tiled format. Multi-resolution tiled images are constructed by first creating multiple copies of the image at different

resolutions. Moreover, at each resolution the image is partitioned into a collection of disjoint tiles.

FLASHPIX®, a registered trademark of the Digital Imaging Group (DIG), is an example of a multi-resolution tiled image format. A FlashPix image is constructed by starting with an original image and recursively sub-sampling it at half the resolution. The recursion continues until the final sub-sampled image is reduced to 64 pixels or less in each dimension. Each resolution level is partitioned into tiles that are 64 x 64 pixels in size. A reference for FLASHPIX® is a document "FlashPix Format Specification," ©1996, 1997, Eastman Kodak Company, the contents of which are hereby incorporated by reference.

Referring to the abovementioned example, for a FlashPix image server to respond with an image at 640 x 480 pixel resolution that contains the full original image which is sub-sampled, is simply a matter of selecting an appropriate close pre-computed resolution. Using the numbers in the example, the successive resolutions are 15,000 x 15,000 pixels, then 7,500 x 7,500 pixels, then 3,750 x 3,750 pixels, then 1,875 x 1,875 pixels, then 937 x 937 pixels, etc. For example, the image server can choose resolution level 937 x 937 pixels and re-sample to 640 x 480 pixels. This is far better than working with a 15,000 x 15,000 pixel image.

FlashPix images are more complicated than simple raster images. The individual 64 x 64 tiles into which each resolution is partitioned are usually JPEG compressed for Internet applications. Furthermore, the FlashPix format specification requires that the tiles be stored in a storage within a Microsoft OLE structured storage file. Structured storage files are compound files composed of multiple storages and streams, where storages are analogous to folders/directories and streams are analogous to files. Although there is overhead in accessing information inside a structured storage file, such files provide a clean interface for a complicated file structure. Structured storage is discussed in Appendix A of the above-referenced FlashPix Format Specification.

For a FlashPix image server to serve an image portion to a client using server-side processing such as the IIP request mentioned above, the following operations are required:

1. Accept the client connection

- i. Wait in accept state for a connection
 - ii. Transmit TCP/IP data
 - iii. Hand off request to a second thread
 - iv. Go back into accept state waiting connection
2. Parse the request
 - i. Check the syntax of the request
 - ii. Format the request into an internal representation
3. Determine the requested tiles
 - i. Determine resolution level
 - ii. Determine portion
 - iii. Determine if re-sampling is necessary
4. Read the tiles from the file
 - i. Open the requested structured storage file
 - ii. Locate the appropriate storage for the previously determined level of resolution
 - iii. Read the meta information (size of tile, location, format, etc.)
 - iv. Read the tile data from a stream
5. Decompress the tiles
 - i. Apply Huffman decoding
 - ii. Apply the Inverse Discrete Cosine Transform
 - iii. De-quantize the coefficients
 - iv. Convert from YUV(4:1:1) to RGB(8:8:8)
6. Stitch the tiles together to form a contiguous image
 - i. Locate positions in memory
 - ii. Copy memory with an appropriate offset
 - iii. Re-size image to desired dimensions
7. Convert (re-compress) the image into the requested format
 - i. Convert from RGB(8:8:8) to YUV(4:1:1)
 - ii. Apply the Discrete Cosine Transform
 - iii. Quantize
 - iv. Apply Huffman encoding
8. Transmit data back to client
 - i. Transmit TCP/IP data

Assignee's co-pending U.S. Patent Application Serial No.

08/ filed November 26, 1997 and entitled A METHOD AND SYSTEM

FOR HTML-DRIVEN INTERACTIVE IMAGE CLIENT, the disclosure of which is hereby incorporated by reference, describes a way to view FlashPix images using the IIP cvt command, without the need for a plug-in or a Java applet. Each interactive user navigation command is implemented through a dynamic HTML page containing a cvt command with appropriate parameters. This puts a maximum strain on the server, as nearly all processing is done on the server side.

The invention described in the aforesaid U.S. Patent Application Serial No. 08/ filed November 26, 1997 and entitled A METHOD AND SYSTEM FOR HTML-DRIVEN INTERACTIVE IMAGE CLIENT utilizes caching to alleviate the computing burden, so that processed data is immediately available for repeated requests from one or many clients for the same image portion. However, when many images are available on the server, each with many image portions that can be viewed, the server must be able to process many demanding cvt commands that each require the extensive processing delineated above, along with many simple requests that can be satisfied immediately with data in cache, or with very little processing and with requisite fairness.

SUMMARY OF THE INVENTION

In U.S. Patent Application Serial No. 08/788,830, filed January 6, 1997 and entitled METHOD AND SYSTEMS FOR SCALABLE REPRESENTATION OF MULTIMEDIA DATA FOR PROGRESSIVE ASYNCHRONOUS TRANSMISSION, the disclosure of which is hereby incorporated by reference, there is described a method and system for delivering on-line interactive media over a server-client network, by means of three databases: (i) a server database, (ii) a client database, and (iii) an interactive database. The mode of delivery described therein is one of partitioning the media into blocks, and pushing the successive blocks from server to client in background, as the media is being interactively accessed by a user. The user has little influence over the data content being delivered to the client database, but has great influence over the interactive database, once the client database is sufficiently large.

The present invention seeks to provide a method and system for enabling digital media to be interactively accessed on-line within a server-client network such as the Internet. In distinction to U.S Patent Application Serial No. 08/788,830, referenced above, which delivers the entire media file to the user, the present invention operates using a "just enough data" policy. That is, rather than providing all of the media from the server computer to the client computer in background, the only data transmitted from the server to the client is the specific data needed by the user; i.e., the specific data necessary to satisfy the interactive request of the user.

There is thus provided in accordance with a preferred embodiment of the present invention a method for communicating media over a network including encoding the media into a server database at a server, downloading from the server database to a client database generally only those portions of the media which are necessary to satisfy user requests, and in response to a user request for a given item of media, determining whether the media is present in the client database, and if not, automatically downloading those portions of the media which are necessary to supply the user with the given item of media from the server database.

Further in accordance with a preferred embodiment of the

present invention the method also includes the step of playing the given item of media to the user.

Still further in accordance with a preferred embodiment of the present invention the method also includes storing in the client database, those portions of the media which are automatically downloaded from the server database, thereby gradually building up the client database.

Moreover in accordance with a preferred embodiment of the present invention the step of playing includes playing a first portion of media received from the client database and playing a second portion of media received from the server database.

Additionally in accordance with a preferred embodiment of the present invention the method also includes identifying which portions of media are needed to satisfy a user request are stored in a client database.

Further in accordance with a preferred embodiment of the present invention the step of identifying includes translating an interactive user request for media into fetch addresses.

Still further in accordance with a preferred embodiment of the present invention the encoding step includes compressing.

Moreover in accordance with a preferred embodiment of the present invention the encoding step includes compressing and also including, prior to the playing step, decompressing the media.

Additionally in accordance with a preferred embodiment of the present invention the identifying step takes place both at the server computer and at a client computer.

Further in accordance with a preferred embodiment of the present invention the media is stored in a directly playable media database prior to playing thereof.

There is also provided in accordance with a preferred embodiment of the present invention apparatus for communicating media over a network including a server which includes a server database having media encoded therein, and a server access controller enabling generally only those portions of the media which are necessary to satisfy user requests to be identified and downloaded from the server database to a client database, and a user computer including a client database, and a client database manager operative to determine whether a user requested item of

media is present in the client database and, if such item is not present, to automatically download those portions of the media which are necessary to supply the user with the user requested item of media from the server database via the server access controller.

Further in accordance with a preferred embodiment of the present invention there is also provided a player for playing the given item of media to the user.

Still further in accordance with a preferred embodiment of the present invention there is also provided apparatus for storing in the client database those portions of the media which are automatically downloaded from the server database, so as to gradually build up the client database.

Moreover in accordance with a preferred embodiment of the present invention the player is operative to play a first portion of media received from the client database and to play a second portion of media received from the server database.

Additionally in accordance with a preferred embodiment of the present invention there is also provided an access controller, to identify which portions of media needed to satisfy a user request are stored on a client database.

Further in accordance with a preferred embodiment of the present invention the access controller is also operative to translate an interactive user request for media into fetch addresses.

Still further in accordance with a preferred embodiment of the present invention there is provided an encoder that performs data compression.

Moreover in accordance with a preferred embodiment of the present invention there is provided an encoder that performs data compression and a decompressor that performs data decompression prior to playing.

Additionally in accordance with a preferred embodiment of the present invention the access controller includes both a server access controller and a client access controller.

Further in accordance with a preferred embodiment of the present invention there is also provided a directly playable media database for storing media prior to playing it.

The present invention also provides a method and system for managing multiple threads to handle requests. As mentioned hereinabove, if requests are processed by a single thread, then when one request is being processed, all subsequent requests have to wait for the first request to finish. If one of the requests takes an excessive amount of time, all subsequent requests are stuck.

The present invention describes a dynamic functionality for thread management, in which additional threads are created and destroyed during processing, depending on how busy the server gets.

As noted above, in the prior art there are two extreme approaches to thread management, neither of which is optimal. The first extreme approach is to restrict to a small number of threads, perhaps even a single thread. This approach has an advantage in that it minimizes overhead involved in operating multiple threads, but it also has a disadvantage in that requests must wait in a queue until the thread is free to accept them. The second extreme approach is to create a large number of threads. This approach has an advantage in that requests are accepted immediately upon arrival, but it also has a disadvantage in that there is great overhead due to operating multiple threads.

Prior art systems for thread management, such as Microsoft IIS and Netscape Enterprise, base their thread management on monitoring client requests. The present invention, in distinction, describes a system that monitors threads.

In a preferred embodiment, the system of the present invention initially creates a small number of threads residing in a thread pool. The threads in the thread pool are monitored for determining how busy they are. Specifically, a special "watchdog" thread is used to monitor the threads and to increment a tick counter associated with each active thread in the thread pool at regular time intervals of 50 msec. A thread's tick counter is set to zero upon start of its activity. Then, at each regular interval when the "watchdog" does its monitoring, the tick counters of the active threads are each incremented by one tick. The tick counters continue to be incremented as long as the threads remain active.

When a request queues up, the system of the present invention does not immediately create a new thread. Rather, the "watchdog"

manages thread creation. Whenever the "watchdog" finds, during its regular check, that the tick counter of a thread has reached 3, it then lowers the priority of this thread and removes it from the thread pool, and creates a new thread to replace it. The old thread that was removed from the thread pool completes its task and dies. Immediately upon its creation, the new thread is free to process a request in the queue. Thus the total number of threads in the thread pool remains a constant.

Since the "watchdog" does its monitoring at regular intervals of 50 msec, and since it removes an active thread from the thread pool when 3 ticks have been registered in the active thread, it turns out that the time during which an active thread can remain in the thread pool is somewhere between 101 msec and 150 msec. Since a new thread is created when an active thread is removed from the thread pool, this is also the maximum time the request queue can build up without any of the queued requests being handled.

Thus the present invention seeks to optimize the number of concurrent active threads and their priorities, in order to handle requests as efficiently as possible while minimizing overhead involved in operating multiple threads.

There is thus provided in accordance with a preferred embodiment of the present invention an image server including at least one processor operative to supply portions of image data to clients in response to multiple requests therefrom, and thread management software operating the at least one processor by causing it to process the requests using at least one of a plurality of threads, the thread management software being characterized in that it initiates a new thread when an existing thread has exceeded a predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the thread management software is characterized in that it initiates a new thread only when an existing thread has exceeded the predetermined metric of busyness.

Additionally in accordance with a preferred embodiment of the present invention the thread management software operates a thread pool containing at least one of the plurality of threads.

Further in accordance with a preferred embodiment of the

present invention the thread management software initiates a new thread when an existing thread in the thread pool has exceeded the predetermined metric of busyness.

Still further in accordance with a preferred embodiment of the present invention the thread management software initiates a new thread only when an existing thread in the thread pool has exceeded the predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the thread management software is operative to open a socket connection for each new incoming request accepted by a thread for processing.

Additionally in accordance with a preferred embodiment of the present invention each existing thread can be in an active state or in a wait state.

Further in accordance with a preferred embodiment of the present invention threads belonging to the thread pool are characterized in that they enter the wait state upon completion of a client request, and are not destroyed.

Still further in accordance with a preferred embodiment of the present invention at least one of the plurality of threads does not belong to the thread pool and is characterized in that it is destroyed upon completion of a client request.

Moreover in accordance with a preferred embodiment of the present invention all of the threads which do not belong to the thread pool are characterized in that they are destroyed upon completion of a client request.

Further in accordance with a preferred embodiment of the present invention the new thread is added to the thread pool.

Additionally in accordance with a preferred embodiment of the present invention if the existing thread belongs to the thread pool, the thread management software is operative to remove the existing thread from the thread pool when the existing thread has exceeded the predetermined metric of busyness.

Further in accordance with a preferred embodiment of the present invention each of the plurality of threads has an associated priority.

Still further in accordance with a preferred embodiment of the present invention the thread management software lowers the priority of the existing thread when the existing thread has exceeded the predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

Additionally in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

Further in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined amount of disk access activity.

Still further in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined amount of memory allocation.

Moreover in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity and at fixed time intervals.

Additionally in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

Further in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

Still further in accordance with a preferred embodiment of the present invention the fixed time intervals are 50 msecs, and the predetermined time duration is the time expended until three ticks have been incremented.

Moreover in accordance with a preferred embodiment of

the present invention the image data is FLASHPIX image data.

Additionally in accordance with a preferred embodiment of the present invention the multiple requests are Internet Imaging Protocol (IIP) requests.

There is also provided in accordance with a preferred embodiment of the present invention a computer thread management system including at least one processor, and thread management software operating the at least one processor by causing it to process requests using at least one of a plurality of threads, the thread management software being characterized in that it initiates a new thread only when an existing thread has exceeded a predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the thread management software is characterized in that it initiates a new thread only when an existing thread has exceeded the predetermined metric of busyness.

Additionally in accordance with a preferred embodiment of the present invention the thread management software operates a thread pool containing at least one of the plurality of threads.

Further in accordance with a preferred embodiment of the present invention the thread management software initiates a new thread when an existing thread in the thread pool has exceeded the predetermined metric of busyness.

Still further in accordance with a preferred embodiment of the present invention the thread management software initiates a new thread only when an existing thread in the thread pool has exceeded the predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the thread management software is operative to open a socket connection for each new incoming request accepted by a thread for processing.

Additionally in accordance with a preferred embodiment of the present invention each existing thread can be in an active state or in a wait state.

Further in accordance with a preferred embodiment of the present invention threads belonging to the thread pool are characterized in

that they enter the wait state upon completion of a client request, and are not destroyed.

Still further in accordance with a preferred embodiment of the present invention at least one of the plurality of threads does not belong to the thread pool and is characterized in that it is destroyed upon completion of a client request.

Moreover in accordance with a preferred embodiment of the present invention all of the threads which do not belong to the thread pool are characterized in that they are destroyed upon completion of a client request.

Further in accordance with a preferred embodiment of the present invention the new thread is added to the thread pool.

Additionally in accordance with a preferred embodiment of the present invention if the existing thread belongs to the thread pool, the thread management software is operative to remove the existing thread from the thread pool when the existing thread has exceeded the predetermined metric of busyness.

Further in accordance with a preferred embodiment of the present invention each of the plurality of threads has an associated priority.

Still further in accordance with a preferred embodiment of the present invention the thread management software lowers the priority of the existing thread when the existing thread has exceeded the predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

Additionally in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

Further in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined amount of disk access activity.

Still further in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined amount of memory allocation.

Moreover in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity and at fixed time intervals.

Additionally in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

Further in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

Still further in accordance with a preferred embodiment of the present invention the fixed time intervals are 50 msec, and the predetermined time duration is the time expended until three ticks have been incremented.

There is also provided in accordance with a preferred embodiment of the present invention an image server method including operating at least one processor for supplying portions of image data to clients in response to multiple requests therefrom including managing processing threads within the at least one processor, thereby causing the processor to process the requests using at least one of a plurality of threads, characterized in that a new thread is initiated when an existing thread has exceeded a predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the managing step initiates a new thread only when an existing thread has exceeded the predetermined metric of busyness.

Additionally in accordance with a preferred embodiment of the present invention the managing step operates a thread pool containing at least one of the plurality of threads.

Further in accordance with a preferred embodiment of the present invention the managing step initiates a new thread when an existing thread in the thread pool has exceeded the predetermined metric of busyness.

Still further in accordance with a preferred embodiment of the present invention the managing step initiates a new thread only when an existing thread in the thread pool has exceeded the predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the managing step includes opening a socket connection for each new incoming request accepted by a thread for processing.

Additionally in accordance with a preferred embodiment of the present invention each existing thread can be in an active state or in a wait state.

Further in accordance with a preferred embodiment of the present invention threads belonging to the thread pool are characterized in that they enter the wait state upon completion of a client request, and are not destroyed.

Still further in accordance with a preferred embodiment of the present invention at least one of the plurality of threads does not belong to the thread pool and is characterized in that it is destroyed upon completion of a client request.

Moreover in accordance with a preferred embodiment of the present invention all of the threads which do not belong to the thread pool are characterized in that they are destroyed upon completion of a client request.

Further in accordance with a preferred embodiment of the present invention the new thread is added to the thread pool.

Additionally in accordance with a preferred embodiment of the present invention if the existing thread belongs to the thread pool, the managing step includes removing the existing thread from the thread pool when the existing thread has exceeded the predetermined metric of busyness.

Further in accordance with a preferred embodiment of the present invention each of the plurality of threads has an associated priority.

Still further in accordance with a preferred embodiment of the present invention the managing step includes lowering the priority of the existing thread when the existing thread has exceeded the predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

Additionally in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

Further in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined amount of disk access activity.

Still further in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined amount of memory allocation.

Moreover in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity and at fixed time intervals.

Additionally in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

Further in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

Still further in accordance with a preferred embodiment of the present invention the fixed time intervals are 50 msecs, and the predetermined time duration is the time expended until three ticks have been incremented.

Moreover in accordance with a preferred embodiment of the present invention the image data is FLASHPIX image data.

Additionally in accordance with a preferred embodiment of the present invention the multiple requests are Internet Imaging Protocol (IIP) requests.

There is also provided in accordance with a preferred embodiment of the present invention a computer thread management method including operating at least one processor, including utilizing thread management software operating the at least one processor by causing it to process requests using at least one of a plurality of threads, the thread management software being characterized in that it initiates a new thread only when an existing thread has exceeded a predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the thread management software is characterized in that it initiates a new thread only when an existing thread has exceeded the predetermined metric of busyness.

Additionally in accordance with a preferred embodiment of the present invention the thread management software operates a thread pool containing at least one of the plurality of threads.

Further in accordance with a preferred embodiment of the present invention the thread management software initiates a new thread when an existing thread in the thread pool has exceeded the predetermined metric of busyness.

Still further in accordance with a preferred embodiment of the present invention the thread management software initiates a new thread only when an existing thread in the thread pool has exceeded the predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the thread management software opens a socket connection for each new incoming request accepted by a thread for processing.

Additionally in accordance with a preferred embodiment of the present invention each existing thread can be in an active state or in a wait state.

Further in accordance with a preferred embodiment of the present invention threads belonging to the thread pool are characterized in that they enter the wait state upon completion of a client request, and are not destroyed.

Still further in accordance with a preferred embodiment of

the present invention at least one of the plurality of threads does not belong to the thread pool and is characterized in that it is destroyed upon completion of a client request.

Moreover in accordance with a preferred embodiment of the present invention all of the threads which do not belong to the thread pool are characterized in that they are destroyed upon completion of a client request.

Further in accordance with a preferred embodiment of the present invention the new thread is added to the thread pool.

Additionally in accordance with a preferred embodiment of the present invention if the existing thread belongs to the thread pool, the thread management software removes the existing thread from the thread pool when the existing thread has exceeded the predetermined metric of busyness.

Further in accordance with a preferred embodiment of the present invention each of the plurality of threads has an associated priority.

Still further in accordance with a preferred embodiment of the present invention the thread management software lowers the priority of the existing thread when the existing thread has exceeded the predetermined metric of busyness.

Moreover in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

Additionally in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

Further in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined amount of disk access activity.

Still further in accordance with a preferred embodiment of the present invention the predetermined metric of busyness is a predetermined amount of memory allocation.

Moreover in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following

initiation of its activity and at fixed time intervals.

Additionally in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

Further in accordance with a preferred embodiment of the present invention the predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

Still further in accordance with a preferred embodiment of the present invention the fixed time intervals are 50 msecs, and the predetermined time duration is the time expended until three ticks have been incremented.

There is also provided in accordance with a preferred embodiment of the present invention a computer thread management system for processing multiple requests by using at least one of a plurality of processing threads including thread monitoring software monitoring the plurality of processing threads, and thread management software managing the plurality of processing threads based upon output from the thread monitoring software.

Moreover in accordance with a preferred embodiment of the present invention the output from the thread monitoring software is a measure of thread busyness.

Additionally in accordance with a preferred embodiment of the present invention the thread management software initiates new threads.

Further in accordance with a preferred embodiment of the present invention the thread management software lowers thread priorities.

Still further in accordance with a preferred embodiment of the present invention the thread management software preserves some of the plurality of threads after processing one of the multiple requests by putting them into a wait state.

Moreover in accordance with a preferred embodiment of the present invention the thread management software allows some of the

plurality of threads to be destroyed after processing one of the multiple requests.

There is also provided in accordance with a preferred embodiment of the present invention a computer thread management method for processing multiple requests using at least one of a plurality of processing threads including monitoring the plurality of processing threads; and managing the plurality of processing threads based upon output from the monitoring step.

Moreover in accordance with a preferred embodiment of the present invention the output from the monitoring step is a measure of thread busyness.

Additionally in accordance with a preferred embodiment of the present invention the managing step initiates new threads.

Further in accordance with a preferred embodiment of the present invention the managing step lowers thread priorities.

Still further in accordance with a preferred embodiment of the present invention the managing step preserves some of the plurality of threads after processing one of the multiple requests by putting them into a wait state.

Moreover in accordance with a preferred embodiment of the present invention the thread management software allows some of the plurality of threads to be destroyed after processing one of the multiple requests.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will be understood and appreciated more fully from the following detailed description, taken in conjunction with the drawings in which:

Fig. 1 is a simplified illustration of a system and method for transmitting digital media data from a server to a client in accordance with a preferred embodiment of the present invention;

Fig. 2 is a simplified flowchart illustrating operation of the system and method of Fig. 1;

Figure 3 is a graph depicting an advantage which the present invention offers over the prior art;

Figure 4 is a screen copy of a performance monitoring graph, explaining behavior corresponding generally to one of the plots in Figure 3 representing the prior art;

Figure 5 is a simplified illustration of a client server system that uses thread management in accordance with a preferred embodiment of the present invention;

Figure 6 is a simplified flowchart of server software within the system of Figure 5;

Figure 7 is a simplified flowchart of a module within the server software of Figure 6 for accepting client requests;

Figure 8 is a simplified flowchart of a module within the server software of Figure 6 for handling client requests; and

Figure 9 is a simplified flowchart of a module within the server software of Figure 6 for managing threads.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Reference is now made to Fig. 1, which is a simplified illustration of a system and method for transmitting digital media data from a server to a client in accordance with a preferred embodiment of the present invention. Digital media is encoded and stored on a server computer 10 within a database 12, for interactive access by users who are networked to the server computer. Interactive access can involve viewing the media, if it is an image, playing it, if it is a movie, or listening to it, if it is an audio signal.

The basic principle of the present invention is for the server to send just enough data to enable the interactive access which the user desires. The data is managed through the use of two databases: the server database 12, residing on the server computer 10, and a client database 14, which is built up on a client computer 16 dynamically, as data is streamed from the server to the client.

Typically the server and client databases contain encoded data, compressed to minimize transmission time. In view of the use of compressed data, there may also optionally be employed a third database 17, as described in U.S. Patent Application Serial No. 08/788,830, referenced above. The third database is an interactive database on the client computer, which contains decoded data, ready for instant access and which serves as a directly playable media database.

For example, suppose the media is a large still image at which the user is gazing, and that the server database contains multiple tiles of the image at different resolutions. When the user initiates a request to gaze at a portion of the image, using a mouse 18 or keyboard 20, the only data transmitted consists of the specific tiles at the specific resolution which are to be displayed. None of the other tiles are transmitted. As the user proceeds to gaze and pan through the image, additional tiles are sent.

The collection of tiles received by the client computer 16 is stored locally in the client computer, in the client database 14. Since the received tiles are thus locally available to the client, whenever the user requests data that was previously accessed, that data can be accessed locally, without the need to transmit it from the server. Since interactive

applications typically involve repeated access to the same data, the client database 14 serves to speed up the interactivity as database 14 is built up.

As another example, suppose the media is a movie that the user is viewing. When the user initiates a request to view a specific frame of the movie, the only data transmitted is that specific frame. None of the other frames are transmitted. As the user advances forward and backward within the movie, additional frames are sent. The frames received by the client computer 16 are stored locally in the client database 14, so that whenever the user returns to a frame which was previously viewed, the data is readily available on the local computer.

Interactive commands input by means of a mouse 18 and keyboard 20 are used to zoom in and out, pan, advance frames, and link to other parts of the media by means of hot spots. Any client computer 16 which is networked to the server can access the media. When a user issues an interactive command, a client access controller 22 identifies the portion of the digital media that is necessary to fulfill the user's request. Initially, before a client database 14 is built up, the client access controller 22 issues a request to a server access controller 24 in the server computer for the data it has identified. Access controller 24 fetches from server database 12 only the portion of data necessary to fulfill the user's request, and then transmits it to the client computer.

The data received by the client computer is then delivered to display 26. In addition, the data is stored in the client database 14. Once the client database is no longer empty, the client access controller 22 first checks whether all of the data it seeks in order to fulfill the user's request is already available in the client database 14. If so, it does not send a request to the server computer, but rather delivers the data from the client database to display 26. Otherwise, it issues a request to the server database 12 via controller 24, as described above, for such data which is not yet available in the client database 14. Once the server transmits such data, the received data is incorporated in the client database 14.

Reference is now made to Fig. 2, which is a simplified flowchart illustrating the operation of the system and method of Fig. 1. At step 30 a user of a client computer initiates an interactive request to access media which is stored on a remote server computer. At step 32 the client

computer identifies what data is necessary in order to satisfy the user's request.

To the extent that the data already resides in the local client database, the data is fetched at step 34 directly from the local computer. To the extent that the data does not already reside in the local client database, at step 36 a request is issued to the server computer to provide the necessary data. At step 38 the server computer fetches the data from its database, and at step 40 the fetched data is sent to the client computer for user access. Finally, at step 42, the received data is added to the client database.

The present invention additionally concerns a method and system for managing multiple threads to handle requests. It describes a way to monitor threads, to use results of the monitoring to manage the threads in an optimal way, and to decide when it is appropriate to initiate a new thread. The monitoring itself can be carried out in its own processing thread. The monitoring system is referred to hereinbelow as a "watchdog."

Prior art systems for thread management, such as those of Microsoft IIS and Netscape Enterprise, base their thread management on monitoring client requests. The present invention, in distinction, describes a system that monitors threads.

In a preferred embodiment, the system of the present invention operates on a thread pool that includes a plurality of threads. These threads can be in an active state or in a wait state. A thread is in an active state when it is busy processing a request. It is in a wait state when it is idle. There can also be additional threads that do not belong to the thread pool. Whereas the threads in the thread pool move into a wait state upon completion of a request, and are not destroyed, threads that do not belong to the thread pool are destroyed upon completion of a request.

The watchdog monitors the status of the threads in the thread pool at regular times, here termed "fixed checkpoint times". Based on specific criteria, the watchdog uses the results of the monitoring, here termed "check results", to decide whether or not to initiate new threads, and, if so, how many new threads to initiate. It also removes certain threads, here termed "slow threads", from the thread pool and lowers their priorities, if the watchdog considers such threads too busy. This is in fair deference to other requests, since a busy thread is usually a reflection of a very demanding

client request.

In a preferred embodiment, the system of the present invention initially creates a small number of threads residing in a thread pool. The threads in the thread pool are monitored in order to determine how busy they are. Specifically, a special watchdog thread is used to monitor the threads and to increment a tick counter associated with each active thread in the thread pool at regular time intervals of 50 msec. A thread's tick counter is set to zero upon the start of its activity. Then, at each regular interval when the watchdog does its monitoring, the tick counters of the active threads are each incremented by one tick. The tick counters continue to be incremented as long as the threads remain active. When a thread in the thread pool completes a request and enters a wait state, its tick counter is reset to zero.

When a request queues up, the system of the present invention does not immediately create a new thread. Rather, the watchdog manages the threads. Whenever the watchdog discovers, during its regular check, that the tick counter of a thread has reached 3, it then lowers the priority of this thread and removes it from the thread pool, and creates a new thread to replace it. The old thread that was removed from the thread pool completes its task and dies. The new thread, immediately upon its creation, is free to process a request in the queue if a queued request exists. Otherwise the thread waits for a request. Thus the total number of threads in the thread pool remains a constant.

Since the watchdog does its checks at regular intervals of 50 msec, and since it removes an active thread from the thread pool when 3 ticks have been registered in the active thread, it turns out that the maximum time during which an active thread can remain in the thread pool is somewhere between 101 msec and 150 msec. The 101 msec corresponds to a thread that started activity at exactly a watchdog checkpoint time, and the 150 msec corresponds to a thread that started its activity 1 msec after a watchdog checkpoint time. Since a new thread is created when an active thread is removed from the thread pool, this maximum time is also the maximum time that a request queue can build up without any of the queued requests being handled.

The use of ticks is believed to be more efficient than having

the watchdog continuously request the system time and compare it to the exact time a thread initiated processing, although the use of system time is more exact.

Thus the present invention seeks to optimize the number of concurrent active threads and their priorities, in order to handle requests as efficiently as possible while minimizing overhead due to operating multiple threads.

Reference is now made to Figure 3, which is a graph depicting an advantage which the present invention offers over the prior art. The graph comes from a server test that was run using requests from simulated clients. Multiple simulated clients each requested access to less than one KByte of data, repeatedly. As the number of multiple clients was varied, the average response time, averaged over 100 repeated requests, for one of the clients was recorded.

The horizontal axis of the graph indicates the number of clients simultaneously making requests, and the vertical axis indicates the average response time, averaged over 100 repeated requests, for a single request (in msec). The curve labeled 102 illustrates the performance of a prior art IIS server, using the thread functionality described hereinabove in the background, wherein a maximum number of threads is pre-set. In the IIS server, multiple threads are created as requests are received, one for each anticipated request. When the maximum limit of allowable concurrent requests/threads is reached, subsequent requests are queued pending availability of a free thread. The IIS server represented by the curve 102 manages its threads by monitoring requests.

A performance monitor was run while the above server test was being conducted, and a screen copy of monitor output for the IIS servers is presented in Figure 4. As can be seen, the IIS server allocated approximately 54 threads. The spikes in the graph correspond to page faults, whereby threads requested pages of memory not readily available. As can be seen in Fig. 4, 15 page faults occurred during the 100 second duration of the graph. This is typical when many threads are active.

The curve labeled 104 in Figure 3 corresponds to a server using a preferred embodiment of the present invention, and is approximately four times faster than the IIS server. This is because the present invention

dynamically allocates threads using a "watchdog" algorithm to monitor threads rather than requests, and was able to process all of the client requests with only 2-3 threads.

The IIS server allocated approximately 54 threads, and assigned threads to each client. These threads competed for memory pages, and as a result page faults were rampant.

In contrast, the server of the present invention, whose performance is indicated in curve 104 waited a short time (approximately 150 msec.) before assigning threads to client requests, to see if an active thread in the thread pool would be freed up and could then be re-used for processing a queued request. Even though the client requests had to wait in a queue, the overall performance was better due to the fact that there were a smaller number of concurrent active threads. Using the IIS server, client requests did not necessarily have to wait in a queue, and were immediately assigned to a waiting thread, but the proliferation of threads caused enough page faults that the overall performance was worse.

Reference is now made to Figure 5, which is a simplified illustration of a client server system that uses thread management in accordance with a preferred embodiment of the present invention. Fig. 5 describes the relationship between a client's request to an image server and the threads managed by the image server to process the request. A client 110 sends a request via IIP, that resides within the HTTP request. This request reaches a server computer 112, which contains server software 114 operative to accept the request and initiate the processing of the request. Processing of the request may entail utilization of a request thread.

During the processing of the request, thread manager 116 concurrently monitors the currently running request threads 118. If a request thread 118 exceeds a maximum limit of processing time allocated per thread, then thread manager 116 performs operations to optimize the performance. These optimizing operations may include creating a new thread and lowering the priority of the thread that exceeded its limit.

Reference is now made to Figure 6, which depicts three main actions that occur concurrently in the server software 114. Each of these actions is enabled by its own set of threads. To accept connections the server software 114 must go into a blocking method, termed `accept()`. In

order that the remainder of the server software 114 will not also be blocked, an Accept Connections step 120 is run in a separate thread. Typically in order to enable an administrator of the image server to stop the server software 114, the Accept Connections step 120 runs the accept method with a timeout of 2 seconds. The accept method is run in a loop, in which a condition for exit from the loop is an administrator request to stop the image server.

Once a connection has been established, the accept method returns a socket object. A socket is described by the four-tuple: remote ip address, remote port, local ip address and local port. The socket object can be used for continued communication with the client.

As seen in Figs. 6, 7 and 8, the socket is passed from the Accept Connections step 120 to a Handle Requests step 122. The Handle Requests step 122 proceeds by reading incoming data sent by a client, parsing the information and processing the request. Since the above processing is likely to take more time than the time between client requests, assuming the existence of multiple clients, it is necessary for the Handle Requests step 122 to return control back to the Accept Connections step 120 which resumes waiting for new client connections. This is done by handing the socket to a separate thread for continued processing. The Handle Requests step 122 thus passes the socket to a thread pool 124.

Due to the nature of threads as described hereinabove, it is desirable to limit the number of concurrent threads while enabling fair handling of concurrent requests. Hence, a Manage Threads step 126 is constantly monitoring the active threads in the thread pool 124. Should an active thread be classified as a "slow thread", i.e. a thread which has exceeded the limit in processing time, the Manage Threads step 126 plays an active role. Otherwise, the Manage Threads step 126 passively watches the thread pool.

It may be appreciated that the criteria for classifying a thread as a "slow thread" need not necessarily be based on processing time. Alternatively or additionally it can be based on any suitable metric, such as a predetermined amount of disk access activity or a predetermined amount of memory allocation.

Reference is now made to Figure 7, which details the

Accept Connections 120 logic. A server socket 132 goes into a blocking Accept step 134, waiting for a client request 130 to arrive. Should a client request 130 not be forthcoming within a timeout limit, typically 2 seconds, the server socket 132 checks if an administrator requested that the server socket 132 exit. Otherwise, the server socket 132 returns to the Accept step 134.

If a client request 130 did occur while the server socket was in Accept step 134, a socket created by the Accept step 134 is passed to a socket queue 136. The socket queue 136 is employed to store client requests 130 pending their processing. As mentioned above, it is not efficient to automatically begin processing each request concurrently, i.e. to start a thread and process the request. Hence, each client request 130 is first placed into the socket queue 136 prior to processing.

Reference is now made to Figure 8, which illustrates a process of handling requests by the Handle Request step 122, which is independent of the Accept Connection step 120. A thread pool 142 stores threads to be used for processing requests. This is more efficient than having each thread die at the end of its processing and restarting new threads, as there is a fixed overhead in creating threads which can be avoided.

The number of threads in the thread pool 142 is preferably proportional to the number of CPUs. A 1:1 relationship is currently recommended. This proportionality limits the inter-dependencies between threads, i.e. the sharing of the CPU resource. The thread pool 142 monitors the activity of the threads in the pool and asks the Free Thread? query 144. If there exists an inactive thread and the socket queue 140 has sockets waiting to be processed, a match is made and the socket is assigned to a thread. The matched thread is marked as active and processing begins in a Process Request step 146. Process Request step 146 includes various different stages mentioned in the Background hereinabove.

Reference is now made to Figure 9, which describes the logic involved in the Manage Threads step 126. A watchdog 152 sleeps for 50 milliseconds, which corresponds to a single tick. This is believed to be more efficient than having the watchdog 152 continuously request the system time and compare it to the exact time a thread initiated processing,

notwithstanding that the use of system time is more exact. At each tick the Thread Active? query 156 is asked of each thread 154 in the thread pool 150. Should no active threads be found, the watchdog 152 returns to a sleep state until the next tick.

When a thread in the thread pool 150 is determined to be active during a watchdog 152 tick by the Thread Active? query 156, then the tick count 158 in the active thread is incremented by one. Next an "Is tick count greater than MAX TICKS?" query 160 is asked. A MAX TICKS value is normally set to 3. Thus if a thread was busy processing for anywhere between 101 and 150 milliseconds, then that thread is labeled a "slow thread."

The labeling of a thread as a slow thread has three consequences enumerated in block 162. First the thread is removed from the thread pool 150, enabling it to die when it finishes processing. Second, the priority of the thread is lowered, reflecting an assumption that difficult requests should not interfere with simple requests in deference to fairness. Third, a new thread is created and placed in the thread pool 150 to replace the slow thread, thus ensuring the presence of a constant number of threads in the thread pool 150 to handle future client requests.

This mechanism ensures that a minimum of threads be used. Only when a request flags the fairness factor does a new thread get created. Furthermore, if requests get blocked attempting to use a shared resource, those threads have their priorities lowered. For example, if many concurrent simple requests are received, they might get blocked on the socket read/write.

It will be appreciated by persons skilled in the art that the present invention is not limited by what has been particularly shown and described hereinabove. Rather the present invention includes combinations and sub-combinations of the various features described hereinabove as well as modifications and extensions thereof which would occur to a person skilled in the art and which do not fall within the prior art.

CLAIMS

1. An image server comprising:
at least one processor operative to supply portions of image data to clients in response to multiple requests therefrom; and
thread management software operating said at least one processor by causing it to process said requests using at least one of a plurality of threads, said thread management software being characterized in that it initiates a new thread when an existing thread has exceeded a predetermined metric of busyness.
2. An image server according to Claim 1 and wherein said thread management software is characterized in that it initiates a new thread only when an existing thread has exceeded said predetermined metric of busyness.
3. An image server according to Claim 1 and wherein said thread management software operates a thread pool containing at least one of said plurality of threads.
4. An image server according to Claim 3 and wherein said thread management software initiates a new thread when an existing thread in said thread pool has exceeded said predetermined metric of busyness.
5. An image server according to Claim 3 and wherein said thread management software initiates a new thread only when an existing thread in said thread pool has exceeded said predetermined metric of busyness.
6. An image server according to Claim 1 and wherein said thread management software is operative to open a socket connection for each new incoming request accepted by a thread for processing.
7. An image server according to Claim 3 and wherein each existing thread can be in an active state or in a wait state.

8. An image server according to Claim 7 and wherein threads belonging to said thread pool are characterized in that they enter said wait state upon completion of a client request, and are not destroyed.
9. An image server according to Claim 3 and wherein at least one of said plurality of threads does not belong to said thread pool and is characterized in that it is destroyed upon completion of a client request.
10. An image server according to Claim 3 and wherein all of the threads which do not belong to said thread pool are characterized in that they are destroyed upon completion of a client request.
11. An image server according to Claim 3 and wherein said new thread is added to said thread pool.
12. An image server according to Claim 3 and wherein if said existing thread belongs to said thread pool, said thread management software is operative to remove said existing thread from said thread pool when said existing thread has exceeded said predetermined metric of busyness.
13. An image server according to Claim 1 and wherein each of said plurality of threads has an associated priority.
14. An image server according to Claim 13 and wherein said thread management software lowers the priority of said existing thread when said existing thread has exceeded said predetermined metric of busyness.
15. An image server according to Claim 1 and wherein said predetermined metric of busyness is a predetermined time duration measured during activity of a thread.
16. An image server according to Claim 4 and wherein said predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

17. An image server according to Claim 1 and wherein said predetermined metric of busyness is a predetermined amount of disk access activity.

18. An image server according to Claim 1 and wherein said predetermined metric of busyness is a predetermined amount of memory allocation.

19. An image server according to Claim 15 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity and at fixed time intervals.

20. An image server according to Claim 15 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

21. An image server according to Claim 16 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

22. An image server according to Claim 15 and wherein said fixed time intervals are 50 msecs, and said predetermined time duration is the time expended until three ticks have been incremented.

23. An image server according to Claim 1 and wherein said image data is FLASHPIX image data.

24. An image server according to Claim 1 and wherein said multiple requests are Internet Imaging Protocol (IIP) requests.

25. A computer thread management system comprising:

at least one processor; and

thread management software operating said at least one processor by causing it to process requests using at least one of a plurality of threads, said thread management software being characterized in that it initiates a new thread only when an existing thread has exceeded a predetermined metric of busyness.

26. A computer thread management system according to Claim 25 and wherein said thread management software is characterized in that it initiates a new thread only when an existing thread has exceeded said predetermined metric of busyness.

27. A computer thread management system according to Claim 25 and wherein said thread management software operates a thread pool containing at least one of said plurality of threads.

28. A computer thread management system according to Claim 27 and wherein said thread management software initiates a new thread when an existing thread in said thread pool has exceeded said predetermined metric of busyness.

29. A computer thread management system according to Claim 27 and wherein said thread management software initiates a new thread only when an existing thread in said thread pool has exceeded said predetermined metric of busyness.

30. A computer thread management system according to Claim 25 and wherein said thread management software is operative to open a socket connection for each new incoming request accepted by a thread for processing.

31. A computer thread management system according to Claim 27 and wherein each existing thread can be in an active state or in a wait state.

32. A computer thread management system according to Claim 31 and wherein threads belonging to said thread pool are characterized in that they enter said wait state upon completion of a client request, and are not destroyed.

33. A computer thread management system according to Claim 27 and wherein at least one of said plurality of threads does not belong to said thread pool and is characterized in that it is destroyed upon completion of a client request.

34. A computer thread management system according to Claim 27 and wherein all of the threads which do not belong to said thread pool are characterized in that they are destroyed upon completion of a client request.

35. A computer thread management system according to Claim 27 and wherein said new thread is added to said thread pool.

36. A computer thread management system according to Claim 27 and wherein if said existing thread belongs to said thread pool, said thread management software is operative to remove said existing thread from said thread pool when said existing thread has exceeded said predetermined metric of busyness.

37. A computer thread management system according to Claim 25 and wherein each of said plurality of threads has an associated priority.

38. A computer thread management system according to Claim 37 and wherein said thread management software lowers the priority of said existing thread when said existing thread has exceeded said predetermined metric of busyness.

39. A computer thread management system according to Claim 25 and wherein said predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

40. A computer thread management system according to Claim 28 and wherein said predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

41. A computer thread management system according to Claim 25 and wherein said predetermined metric of busyness is a predetermined amount of disk access activity.

42. A computer thread management system according to Claim 25 and wherein said predetermined metric of busyness is a predetermined amount of memory allocation.

43. A computer thread management system according to Claim 39 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity and at fixed time intervals.

44. A computer thread management system according to Claim 39 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

45. A computer thread management system according to Claim 40 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

46. A computer thread management system according to Claim 39 and wherein said fixed time intervals are 50 msecs, and said predetermined time duration is the time expended until three ticks have been incremented.

47. An image server method comprising:

operating at least one processor for supplying portions of image data to clients in response to multiple requests therefrom including:

managing processing threads within said at least one processor, thereby causing said processor to process said requests using at least one of a plurality of threads, characterized in that a new thread is initiated when an existing thread has exceeded a predetermined metric of busyness.

48. A method according to Claim 47 and wherein said managing step initiates a new thread only when an existing thread has exceeded said predetermined metric of busyness.

49. A method according to Claim 47 and wherein said managing step operates a thread pool containing at least one of said plurality of threads.

50. A method according to Claim 49 and wherein said managing step initiates a new thread when an existing thread in said thread pool has exceeded said predetermined metric of busyness.

51. A method according to Claim 49 and wherein said managing step initiates a new thread only when an existing thread in said thread pool has exceeded said predetermined metric of busyness.

52. A method according to Claim 47 and wherein said managing step includes opening a socket connection for each new incoming request accepted by a thread for processing.

53. A method according to Claim 49 and wherein each existing thread can be in an active state or in a wait state.

54. A method according to Claim 53 and wherein threads belonging to said thread pool are characterized in that they enter said wait state upon completion of a client request, and are not destroyed.

55. A method according to Claim 49 and wherein at least one of said plurality of threads does not belong to said thread pool and is characterized in that it is destroyed upon completion of a client request.

56. A method according to Claim 49 and wherein all of the threads which do not belong to said thread pool are characterized in that they are destroyed upon completion of a client request.

57. A method according to Claim 49 and wherein said new thread is added to said thread pool.

58. A method according to Claim 49 and wherein if said existing thread belongs to said thread pool, said managing step includes removing said existing thread from said thread pool when said existing thread has exceeded said predetermined metric of busyness.

59. A method according to Claim 47 and wherein each of said plurality of threads has an associated priority.

60. A method according to Claim 59 and wherein said managing step includes lowering the priority of said existing thread when said existing thread has exceeded said predetermined metric of busyness.

61. A method according to Claim 47 and wherein said predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

62. A method according to Claim 50 and wherein said predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

63. A method according to Claim 47 and wherein said predetermined metric of busyness is a predetermined amount of disk access activity.

64. A method according to Claim 47 and wherein said predetermined metric of busyness is a predetermined amount of memory allocation.

65. A method according to Claim 61 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity and at fixed time intervals.

66. A method according to Claim 61 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

67. A method according to Claim 62 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

68. A method according to Claim 61 and wherein said fixed time intervals are 50 msecs, and said predetermined time duration is the time expended until three ticks have been incremented.

69. A method according to Claim 47 and wherein said image data is FLASHPIX image data.

70. A method according to Claim 47 and wherein said multiple requests are Internet Imaging Protocol (IIP) requests.

71. A computer thread management method comprising:
operating at least one processor, including:
utilizing thread management software
operating said at least one processor by causing it to process requests using
at least one of a plurality of threads, said thread management software being
characterized in that it initiates a new thread only when an existing thread

has exceeded a predetermined metric of busyness.

72. A method according to Claim 71 and wherein said thread management software is characterized in that it initiates a new thread only when an existing thread has exceeded said predetermined metric of busyness.

73. A method according to Claim 71 and wherein said thread management software operates a thread pool containing at least one of said plurality of threads.

74. A method according to Claim 73 and wherein said thread management software initiates a new thread when an existing thread in said thread pool has exceeded said predetermined metric of busyness.

75. A method according to Claim 73 and wherein said thread management software initiates a new thread only when an existing thread in said thread pool has exceeded said predetermined metric of busyness.

76. A method according to Claim 71 and wherein said thread management software opens a socket connection for each new incoming request accepted by a thread for processing.

77. A method according to Claim 73 and wherein each existing thread can be in an active state or in a wait state.

78. A method according to Claim 77 and wherein threads belonging to said thread pool are characterized in that they enter said wait state upon completion of a client request, and are not destroyed.

79. A method according to Claim 73 and wherein at least one of said plurality of threads does not belong to said thread pool and is characterized in that it is destroyed upon completion of a client request.

80. A method according to Claim 73 and wherein all of the

threads which do not belong to said thread pool are characterized in that they are destroyed upon completion of a client request.

81. An image server according to Claim 73 and wherein said new thread is added to said thread pool.

82. A method according to Claim 73 and wherein if said existing thread belongs to said thread pool, said thread management software removes said existing thread from said thread pool when said existing thread has exceeded said predetermined metric of busyness.

83. A method according to Claim 71 and wherein each of said plurality of threads has an associated priority.

84. A method according to Claim 83 and wherein said thread management software lowers the priority of said existing thread when said existing thread has exceeded said predetermined metric of busyness.

85. A method according to Claim 71 and wherein said predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

86. A method according to Claim 74 and wherein said predetermined metric of busyness is a predetermined time duration measured during activity of a thread.

87. A method according to Claim 71 and wherein said predetermined metric of busyness is a predetermined amount of disk access activity.

88. A method according to Claim 71 and wherein said predetermined metric of busyness is a predetermined amount of memory allocation.

89. A method according to Claim 85 and wherein said

predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity and at fixed time intervals.

90. A method according to Claim 85 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

91. A method according to Claim 86 and wherein said predetermined time duration is determined by incrementing ticks for each thread which is in an active state following initiation of its activity at fixed time intervals which are not measured from the onset of its activity.

92. A method according to Claim 85 and wherein said fixed time intervals are 50 msecs, and said predetermined time duration is the time expended until three ticks have been incremented.

93. A computer thread management system for processing multiple requests by using at least one of a plurality of processing threads comprising:

thread monitoring software monitoring said plurality of processing threads; and

thread management software managing said plurality of processing threads based upon output from said thread monitoring software.

94. A computer thread management system according to Claim 93 and wherein said output from said thread monitoring software is a measure of thread busyness.

95. A computer thread management system according to Claim 93 and wherein said thread management software initiates new threads.

96. A computer thread management system according to Claim 93 and wherein said thread management software lowers thread priorities.

97. A computer thread management system according to Claim 93 and wherein said thread management software preserves some of said plurality of threads after processing one of said multiple requests by putting them into a wait state.

98. A computer thread management system according to Claim 93 and wherein said thread management software allows some of said plurality of threads to be destroyed after processing one of said multiple requests.

99. A computer thread management method for processing multiple requests using at least one of a plurality of processing threads comprising:

monitoring said plurality of processing threads; and
managing said plurality of processing threads based upon
output from said monitoring step.

100. A method according to Claim 99 and wherein said output from said monitoring step is a measure of thread busyness.

101. A method according to Claim 99 and wherein said managing step initiates new threads.

102. A method according to Claim 99 and wherein said managing step lowers thread priorities.

103. A method according to Claim 99 and wherein said managing step preserves some of said plurality of threads after processing one of said multiple requests by putting them into a wait state.

104. A method according to Claim 99 and wherein said thread management software allows some of said plurality of threads to be destroyed after processing one of said multiple requests.

105. A method for communicating media over a network comprising the steps of:

encoding the media into a server database at a server;

downloading from the server database to a client database generally only those portions of the media which are necessary to satisfy user requests; and

in response to a user request for a given item of media, determining whether said media is present in the client database, and if not, automatically downloading those portions of the media which are necessary to supply the user with the given item of media from the server database.

106. A method according to claim 105 and also comprising the step of playing the given item of media to the user.

107. A method according to claim 106 and also comprising the step of storing in said client database, those portions of the media which are automatically downloaded from the server database, thereby to gradually build up the client database.

108. A method according to claim 106 and wherein the step of playing includes playing a first portion of media received from the client database and playing a second portion of media received from the server database.

109. A method according to claim 105 and also comprising the step of identifying which portions of media needed to satisfy a user request are stored on a client database.

110. A method according to claim 109 and wherein the step of identifying includes the step of translating an interactive user request for media into fetch addresses.

111. A method according to claim 105 and wherein said encoding step includes the step of compressing.

112. A method according to claim 106 wherein said encoding step includes the step of compressing and also comprising, prior to said playing step, the step of decompressing said media.

113. A method according to claim 110 and wherein said identifying step takes place both at the server computer and at a client computer.

114. A method according to claim 106 and wherein said media is stored in a directly playable media database prior to playing thereof.

115. Apparatus for communicating media over a network comprising:

a server including:

a server database having media encoded

therein; and

a server access controller enabling generally only those portions of the media which are necessary to satisfy user requests to be identified and downloaded from the server database to a client database; and

a client computer operated by a user including:

a client database; and

a client database manager operative to determine whether a user requested item of media is present in the client database and, if such item is not present, to automatically download those portions of the media which are necessary to supply the user with the user requested item of media from the server database via the server access controller.

116. Apparatus according to claim 115 and also comprising a player for playing the given item of media to the user.

117. Apparatus according to claim 116 and also comprising apparatus for storing in said client database those portions of the media which are automatically downloaded from the server database, thereby to

gradually build up the client database.

118. Apparatus according to claim 116 and wherein the player is operative to play a first portion of media received from the client database and to play a second portion of media received from the server database.

119. Apparatus according to claim 115 and also comprising an access controller, identifying which portions of media needed to satisfy a user request are stored on a client database.

120. Apparatus according to claim 119 and wherein the access controller is also operative to translate an interactive user request for media into fetch addresses.

121. Apparatus according to claim 115 and also comprising an encoder that performs data compression.

122. Apparatus according to claim 116 and also comprising an encoder that performs data compression and a decompressor that performs data decompression prior to playing.

123. Apparatus according to claim 120 and wherein said access controller includes both a server access controller and a client access controller.

124. Apparatus according to claim 116 and also comprising a directly playable media database for storing media prior to playing thereof.

1/9




FIGURE 1





FIGURE 2


3/9

FIG. 3



4/9

FIG. 4
PRIOR ART



5/9

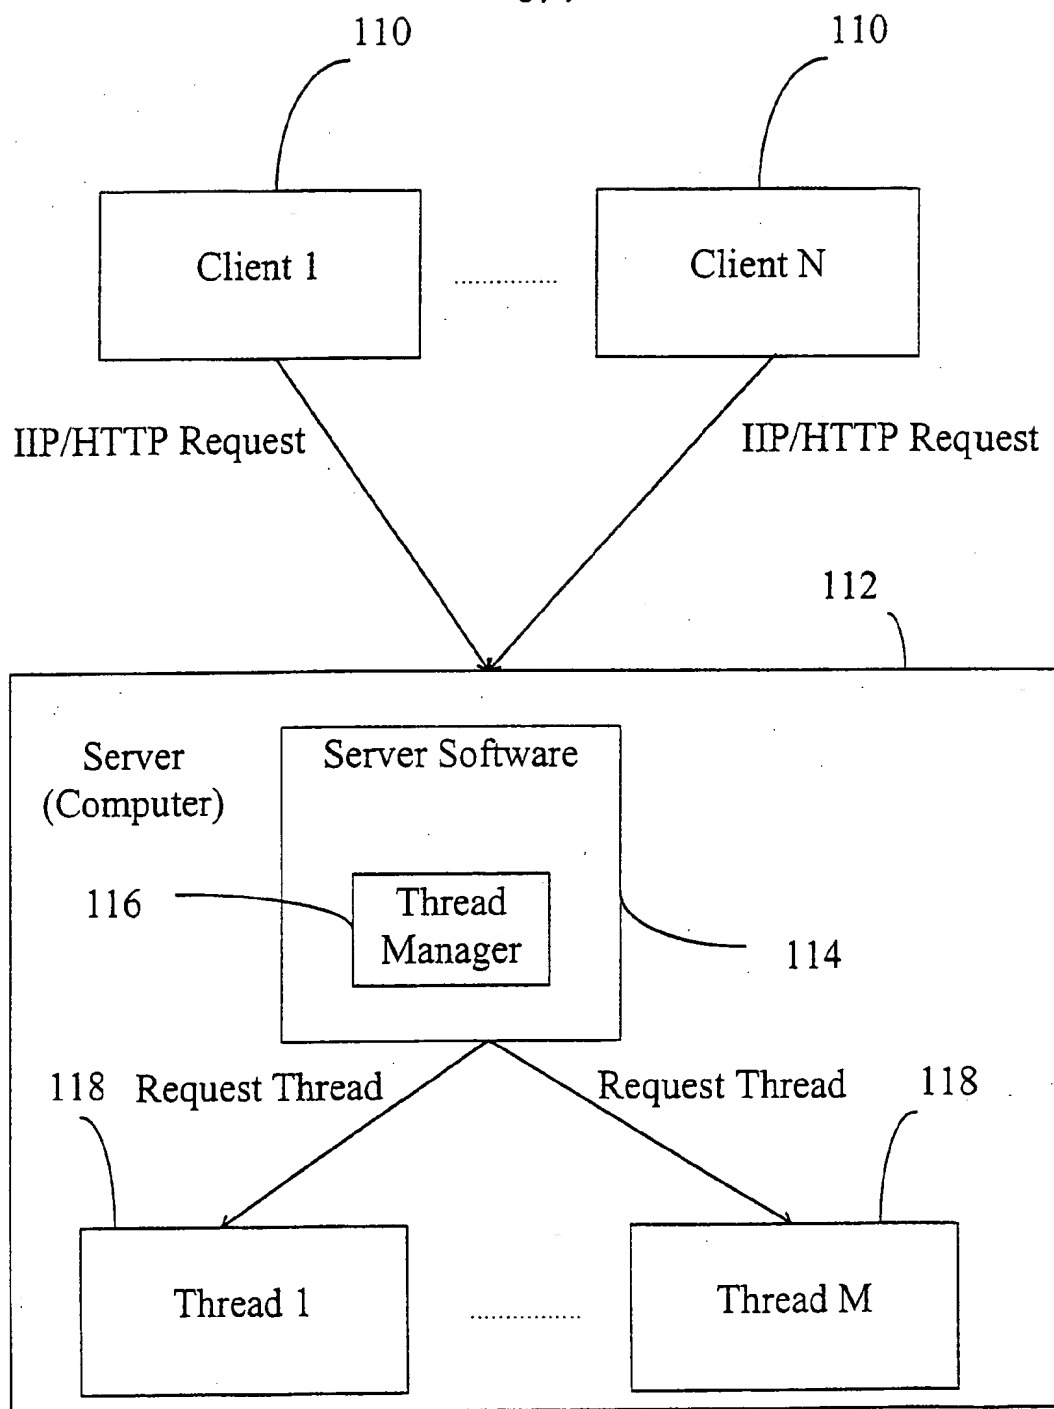


Figure 5

6/9

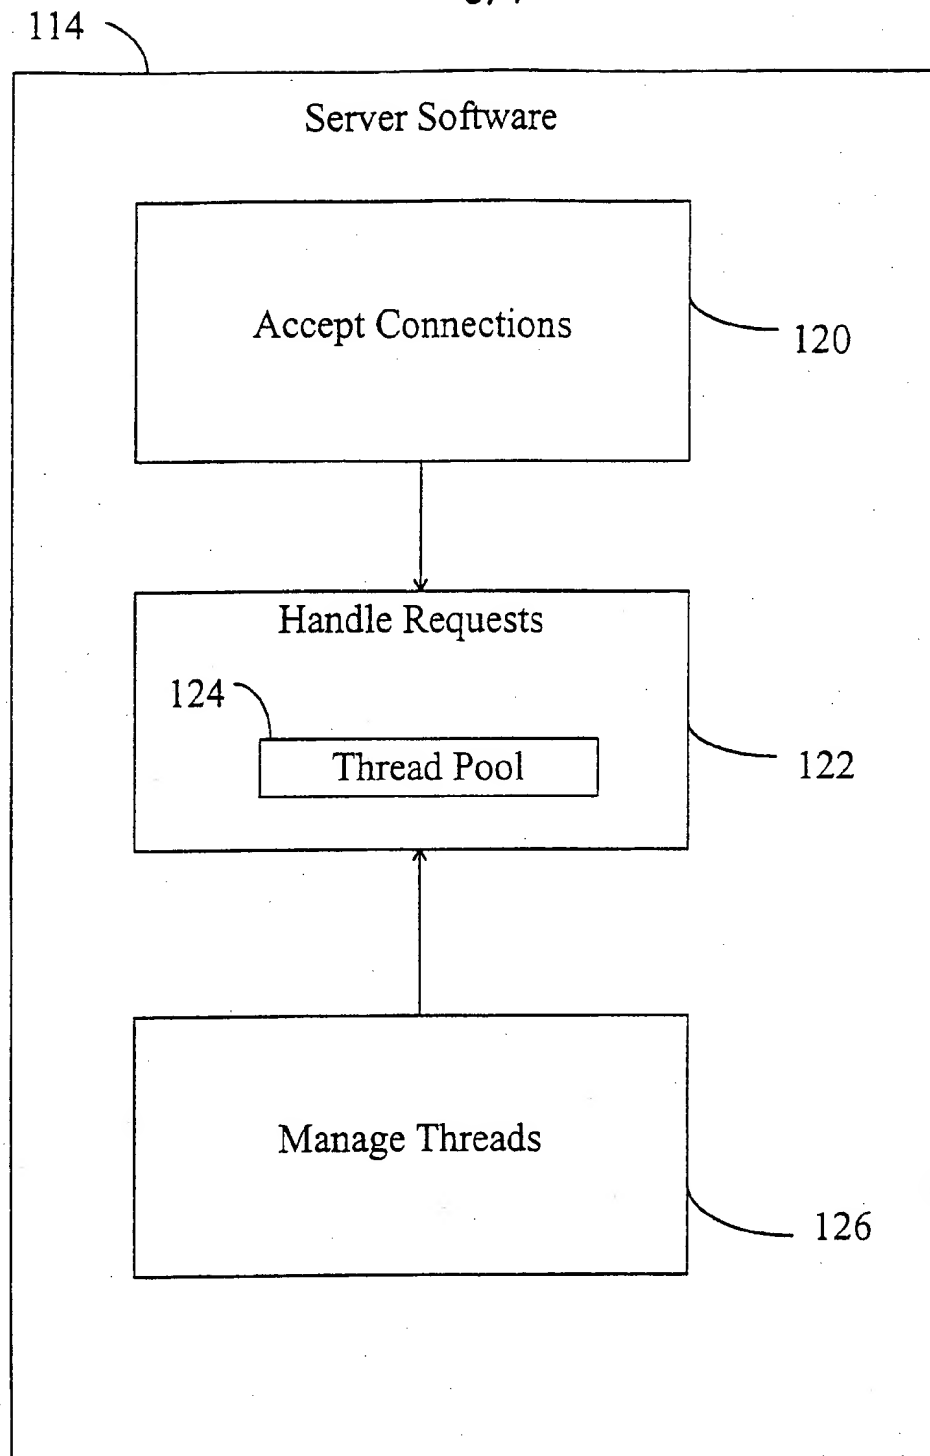


Figure 6

7/9

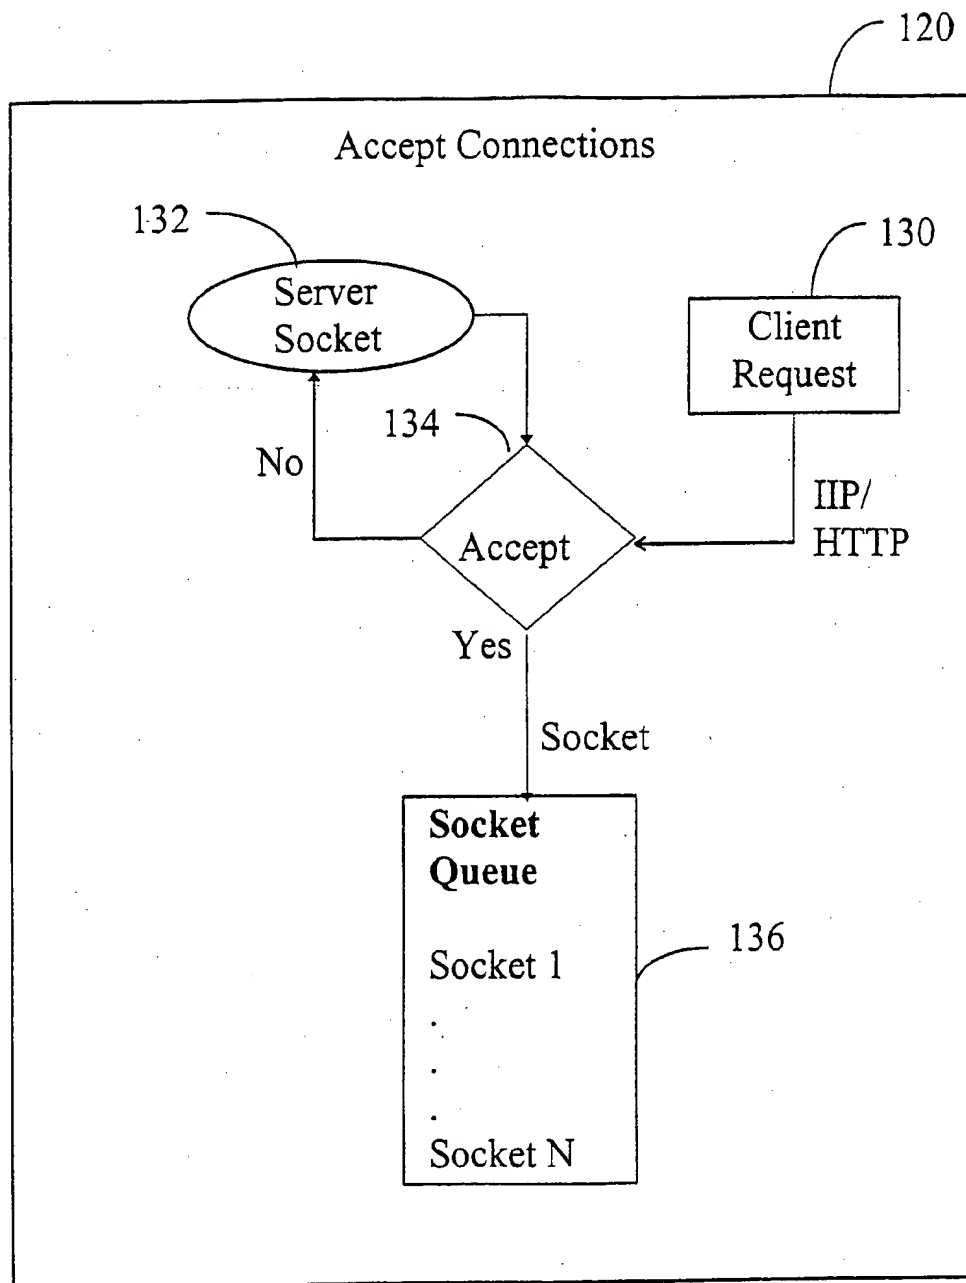


Figure 7

8/9

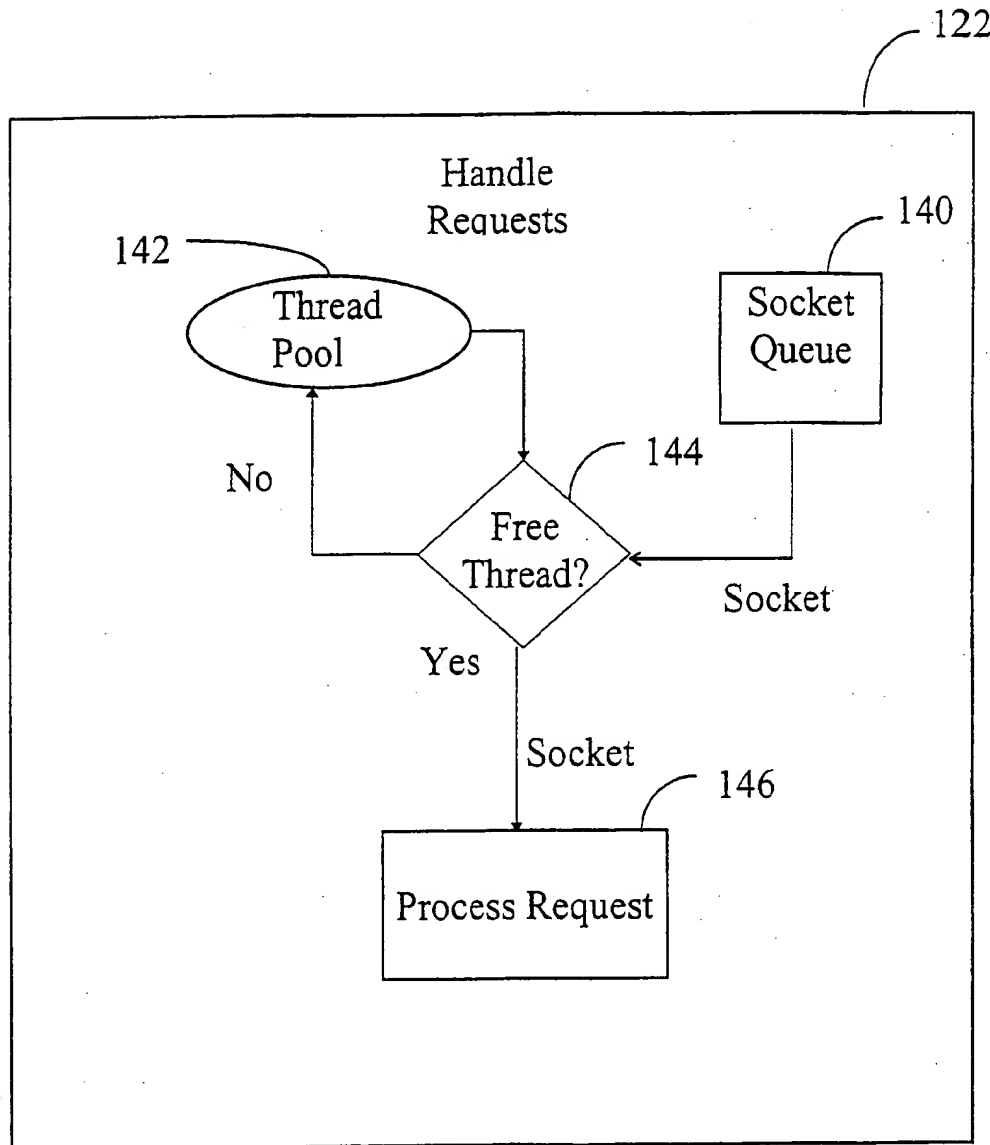


Figure 8

9/9

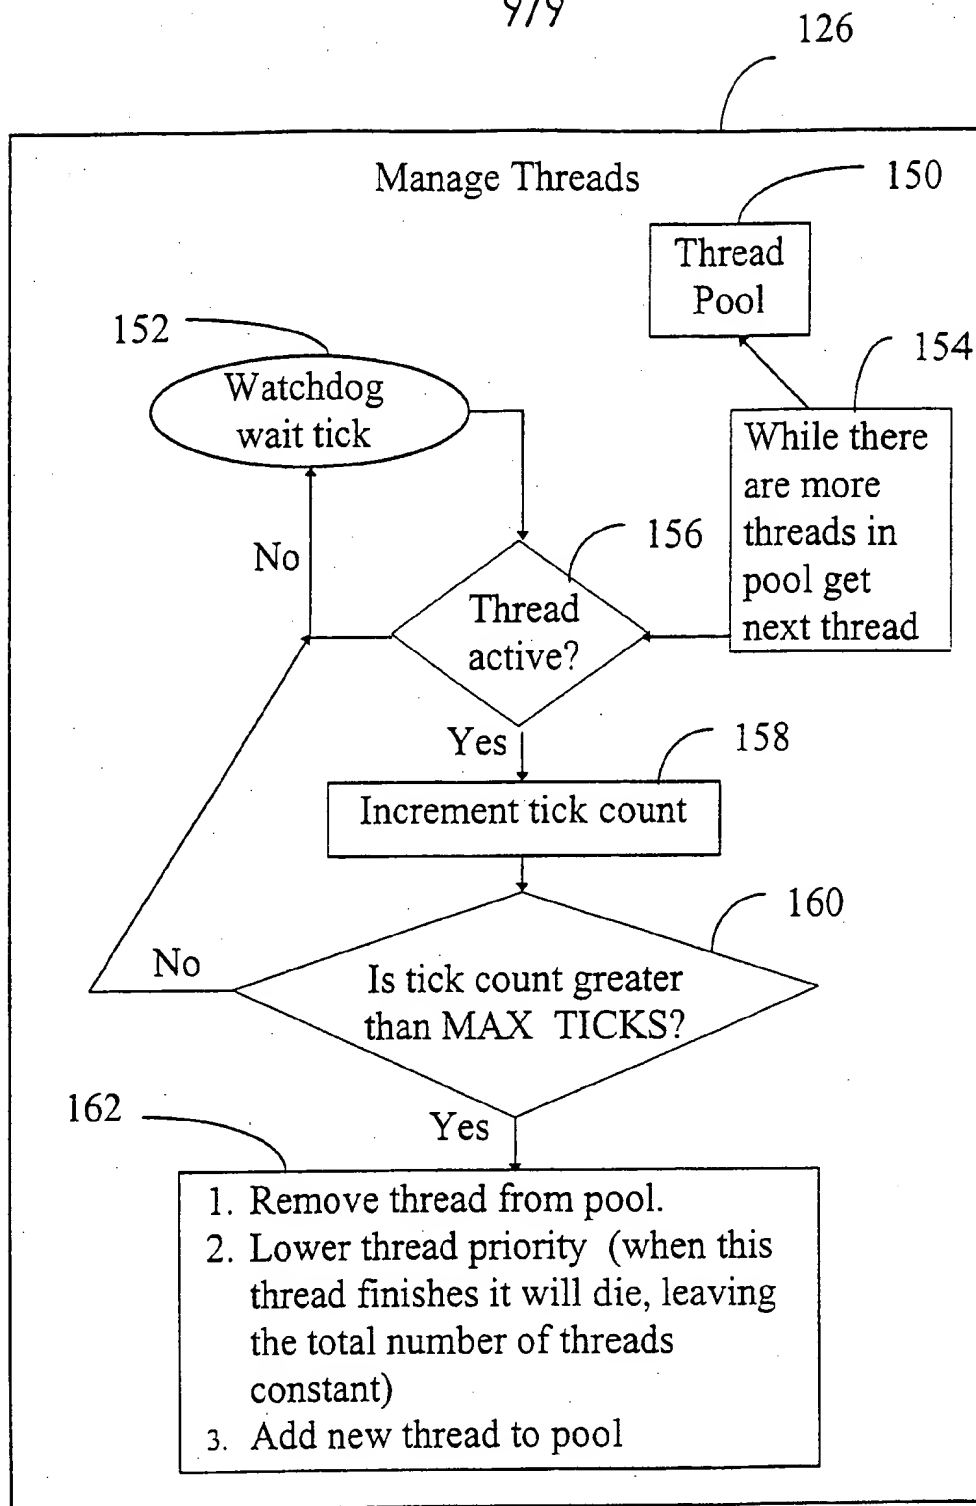


Figure 9

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- ☐ BLACK BORDERS
- ☐ IMAGE CUT OFF AT TOP, BOTTOM OR SIDES
- ☒ FADED TEXT OR DRAWING
- ☐ BLURRED OR ILLEGIBLE TEXT OR DRAWING
- ☐ SKEWED/SLANTED IMAGES
- ☐ COLOR OR BLACK AND WHITE PHOTOGRAPHS
- ☐ GRAY SCALE DOCUMENTS
- ☐ LINES OR MARKS ON ORIGINAL DOCUMENT
- ☐ REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY
- ☐ OTHER: _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.